

Amendments to the Drawings:

Attached are copies of Figures 3 and 4, marked in red to show proposed corrections, and substitute drawing sheets of Figures 3 and 4 incorporating the corrections. Specifically, in decision block 32 of Figure 3, the legend has been changed to $-i \geq 2^n$. A similar correction has been made to decision block 42 in Figure 4.

REMARKS

The specification has been amended at pages 4 and 5 as required by the Examiner. In addition, the equation on page 8, line 2, has been corrected to eliminate the extra term " $w_{1,1}$ ", which appeared twice in the right-hand matrix so that the matrix has only six entries to match the number of the columns in the middle matrix.

Figures 3 and 4 of the drawings have been amended as required by the Examiner. Specifically, in decision block 32 of Figure 3, the legend has been changed to $-i \geq 2^n -$. A similar correction has been made to decision block 42 in Figure 4.

Claims 1 and 3 to 12 appear in the application. The indication that claims 3 and 8 to 10 are directed to allowable subject matter is noted with appreciation. By this amendment, claim 3 has been amended to be in independent form by incorporating the limitations of claim 2, and claim 2 has been canceled. With this amendment, claim 3 is in condition for immediate allowance.

The present invention relates to constructing codes for which the encoding and correcting algorithms can be executed fast or "on-the-fly" and has particular application in reading data from arrays of hard disks in data processing systems, commonly called RAID arrays for Redundant Array of Inexpensive (or Independent) Disks. Failure correction codes for RAID systems which correct more than one failed disk use coefficients for the parity equations. These coefficients, which cannot be all "1s" or "0s", are comprised of several bits and, during the encoding and correcting stages, these coefficients multiply the data which is broken into chunks whose size is the same as that of the coefficients. Even when those coefficients are 0-1 matrices as in Ghosh et al., U.S. Patent No. 6,823,425, the data is broken into chunks whose size is the dimension of the matrices so that it can be multiplied by the matrices. It is this multiplication of the coefficients with the data which is time consuming when implemented in software. First, the appropriate chunks of data, for example bytes, have to be extracted from the word used by the microprocessor, for example the bytes extracted from the word, and then the multiplication by the coefficients have to be

performed. The results of these multiplications are then exclusive ORed (XORed). When implemented in software, the multiplication operation is usually done by table lookup so as to avoid the time consuming operations on the bits. That is why RAID systems often use an ASIC (Application Specific Integrated Circuit) to perform the multiplication of the data by the coefficients, be they elements of a finite field or matrices.

In contrast with the prior art, the disclosed and claimed invention describes a RAID error correction scheme for correcting more than one failed disk which does not call for multiplication by coefficients which are not 0 or 1. Thus, the invention can be implemented in software using an off-the-shelf microprocessor, using only the XOR operation on words, without needing to extract parts of the word, resorting to table lookup, performing operations on bits, or needing the assistance of an ASIC. In the sentence above, the term "word" is used to be that which the microprocessor, rather than the code, determines. That is, the term "word" as used in the disclosed and claimed invention is used to denote a machine word. Thus, if the microprocessor is a 32-bit microprocessor, then a word is 32 bits long, and if it is a 64-bit microprocessor, then a word is 64 bits long. This is an accepted definition of "word" as indicated by the entry on page 418 of Microsoft Press' *Computer Dictionary*, Second Edition (1994), a copy of which is attached.

Claim 1 was rejected under 35 U.S.C. §102(e) as being anticipated by U.S. Patent No. 6,694,479 to Murthy et al. This rejection is respectfully traversed for the reason that Murthy et al. do not show, suggest or otherwise teach the claimed invention.

The Examiner cites equations (1), (2) and (3) in columns 6 and 7 in the patent to Murthy et al. However, these equations do use multiplications by the coefficients of the parity equations. Equation (2), in column 6, includes terms such as $3d_3$ and $11d_{11}$, and equation (3) includes terms such as 5^2d_5 and 7^2d_7 . It is precisely these type of multiplications of the data by the coefficients of the parity equations, such as $3d_3$ and 7^2d_7 , which claimed inventions avoids. Claim 1 recites the step of "performing **only** exclusive OR operations on words for error correcting codes with four or more check symbols which can correct as many

errors as there are check symbols” (emphasis added). In fact, Murthy et al. explain at column 3, line 55, to column 4, line 13, that when the code is defined by two or more parity equations (that is, more than one failure can be corrected), some of the coefficients have to be other than 0 or 1, and therefore the multiplication by these coefficients is a nontrivial multiplication. Moreover, as is clear from the content of the patent to Murthy et al., the term “word” is used to have a size such that the multiplication by the coefficients can be performed, and thus the term “word” is not used in the sense of the disclosed invention. For the reasons above, it is submitted that claim 1 is not anticipated by Murthy et al. and, moreover, claim 1 is not made obvious under 35 U.S.C. §103 by the patent to Murthy et al. since there is no teaching of correcting errors by performing only exclusive OR operations.

Claims 2, 4 to 7, 11, and 12 were rejected under 35 U.S.C. §102(e) as being anticipated by U.S. Patent No. 6,823,425 to Ghosh et al. This rejection is moot as to claim 2 which has been canceled with incorporation of the limitations of claim 2 into claim 3. As to the remaining claims 4 to 7, 11 and 12, the rejection is respectfully traversed for the reason that Ghosh et al. do not show, suggest or otherwise teach the claimed invention.

The Examiner takes the position that Ghosh et al. teach “encoding and correcting algorithms [which] involve only exclusive OR (XOR) operations of words”; however, Applicant disagrees. Ghosh et al. use the term “word” for that which their matrices can multiply, i.e., having the same number of bits as the dimension of the matrices. The set of matrices which Ghosh et al. generate operate on bits, not the words of the microprocessor. As Ghosh et al. state in Column 9, lines 9–13: “The following representation shows explicitly which **bits** of word w must be XORed in order to obtain its translation $T_k(w)$. For example, if the i^{th} row is 00101001, then the i^{th} bit of the translated word is obtained by XORing together the third, fifth and eighth **bit** of the input word.” (emphasis added). The disadvantages of using bit operations was described above, and it is precisely that which the claimed invention avoids, and where it differs from prior art. The fact that only XOR words of the microprocessor are used is where the claimed invention improves on the prior art, and is the crucial point for which patentable weight should be given.

The claimed invention is about performing only XORs of words, avoiding bit operations (or table lookups) altogether. Thus, claim 4 recites, *inter alia*, “encoding data and correcting erasure errors using **only** exclusive OR operations on words” (emphasis added). Claim 5 is dependent on claim 4 and, therefore, incorporates the same limitation. Claim 6 recites, *inter alia*, “reading data from main volatile memory and encoding the data using **only** XOR operations to generate a correction code” and “reconstructing erasure errors in the data read from the auxiliary array of non-volatile storage devices using **only** XOR operations to generate reconstructed data” (emphasis added). Claim 7, which is dependent on claim 6, adds that “the code whose encoding and correcting algorithms involve **only** XOR operations of words is a (3,3) code of distance four” (emphasis added). Claim 11 recites, *inter alia*, that the “encoding and correcting algorithms involve **only** exclusive OR (XOR) operations of words, data read from said main volatile memory being encoded by said encoding means using **only** XOR operations to generate a correcting code” and erasure errors are reconstructed “in the data read from the auxiliary array of non-volatile memory devices using **only** XOR operations to generate reconstructed data” (emphasis added).

The invention shows how to construct RAID failure recovery that, when implemented on an off-the-shelf microprocessor uses only XORs of words (and loads and stores), and does not need the assistance of an ASIC, nor does it need to perform bit operations or table-lookup. It is important to be able to use an off-the-shelf microprocessor without having to add an ASIC for cost reasons, and it is important to avoid having to do table look-ups or bit operations in software because they are time consuming and slow down the failure recovery.

The prior art made of record but not relied upon has been reviewed; however, none of the prior art show, suggest or otherwise teach using only XORs of words for RAID failure recovery.

In view of the foregoing, it is respectfully requested that the application be reconsidered, that claims 1, 4 to 7, 11 and 12 be allowed together with claims 3 and 8 to 10, and that the application be passed to issue.

Should the Examiner find the application to be other than in condition for

allowance, the Examiner is requested to contact the undersigned at the local telephone number listed below to discuss any other changes deemed necessary in a telephonic or personal interview.

A provisional petition is hereby made for any extension of time necessary for the continued pendency during the life of this application. Please charge any fees for such provisional petition and any deficiencies in fees and credit any overpayment of fees to Attorney's Deposit Account No. 50-2041.

Respectfully submitted,

A handwritten signature in black ink, appearing to read 'C. Lamont Whitham', is written over a horizontal line.

C. Lamont Whitham
Reg. No. 22,424

Whitham, Curtis & Christofferson, P.C.
11491 Sunset Hills Road, Suite 340
Reston, VA 20190

Tel. (703) 787-9400
Fax. (703) 787-7557

Customer No.: 30743

MICROSOFT PRESS®

COMPUTER DICTIONARY

SECOND EDITION



THE COMPREHENSIVE
STANDARD FOR
BUSINESS, SCHOOL,
LIBRARY, AND HOME



PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 1994 by Microsoft Press

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data

Microsoft Press computer dictionary : the comprehensive standard for
business, school, library, and home / Microsoft Press. -- 2nd ed.

p. cm.

ISBN 1-55615-597-2

1. Computers--Dictionaries. 2. Microcomputers--Dictionaries.

I. Microsoft Press. II. Title: Computer dictionary.

QA76.15.M54 1993

004'.03--dc20

93-29868

CIP

Printed and bound in the United States of America.

1 2 3 4 5 6 7 8 9 MLML 9 8 7 6 5 4

Distributed to the book trade in Canada by Macmillan of Canada, a division of Canada
Publishing Corporation.

Distributed to the book trade outside the United States and Canada by
Penguin Books Ltd.

Penguin Books Ltd., Harmondsworth, Middlesex, England
Penguin Books Australia Ltd., Ringwood, Victoria, Australia
Penguin Books N.Z. Ltd., 182-190 Wairau Road, Auckland 10, New Zealand

British Cataloging-in-Publication Data available.

Project Editor: Casey D. Doyle

Manuscript Editor: Alice Copp Smith

Technical Editors: Mary DeJong, Jeff Carey, Dail Magee, Jr., Jim Fuchs, Seth McEvoy



dows NT Advanced Server also offers advanced hard disk fault-tolerance features, such as mirroring, and additional connectivity. *See also* Windows NT.

wire-frame model In computer graphics applications such as CAD programs, a display of a three-dimensional object composed of separate lines that resemble strands of wire joined to create a model. *Compare* solid model.

wire-pin printer *See* dot-matrix printer.

wire-wrapped circuits Circuits constructed on perforated boards using wire instead of the metal traces found on printed circuit boards. The bare ends of insulated wires are wrapped around the long pins of special wire-wrap integrated circuit sockets. Wire-wrapped circuits are generally handmade one-of-a-kind devices used for prototyping and research in electrical engineering. Their advantage is that the wires are easily unwrapped and the pin-to-pin connections changed, permitting circuit designers to experiment with a circuit's design without having to lay out and etch a new printed circuit board. *Compare* printed circuit board.

wizard Someone who is adept at making computers perform their "magic"; an outstanding and creative programmer or a power user.

wizzywig *See* WYSIWYG.

word The native unit of storage on a particular machine. Depending on the microprocessor, a word can be an 8-bit, a 16-bit, or a 32-bit quantity.

word-addressable processor A processor that cannot access an individual byte of memory but can access only a larger unit. In order to perform operations on an individual byte, the processor must read and write memory in the larger unit. For example, a word-addressable processor might read a word (two bytes) from memory at one time, add a value to only one of the bytes, and then write the word back to memory.

word length Typically, a standard data unit (8-bit, 16-bit, and 32-bit words are the most common) in a particular computer, representing both the largest amount of data that can be handled by the microprocessor in one operation and also, as a rule, the width of the main data bus (the hard-

ware pathway that carries information from place to place within the computer).

word processing Abbreviated WP. The act of entering text and editing with a word processor. *See also* word processor.

word processor An application program for manipulating text-based documents; the electronic equivalent of paper, pen, typewriter, eraser, and, most likely, dictionary and thesaurus. Word processors run the gamut from simple through complex, but all ease the tasks associated with editing documents (deleting, inserting, rewording, and so on). Depending on the program and the equipment in use, word processors can display documents either in text mode, using highlighting, underlining, or color to represent italics, boldfacing, and other such formatting, or in graphics mode, wherein formatting and, sometimes, a variety of fonts appear on the screen as they will on the printed page. All word processors offer at least limited facilities for document formatting, such as font changes, page layout, paragraph indentation, and the like. Some word processors can also check spelling, find synonyms, incorporate graphics created with another program, correctly align mathematical formulas, create and print form letters, perform calculations, display documents in multiple on-screen windows, and enable users to record macros that simplify difficult or repetitive operations. *Compare* editor, line editor.

wordwrap The ability of a word-processing program to break lines of text automatically to stay within the page margins of a document. Line breaks created by wordwrap are known as soft returns. *See also* hard return, soft return.

worksheet A term used to describe a data file created by and used with an electronic spreadsheet program. Also, an alternative name for a spreadsheet. *See also* spreadsheet program.

workstation In general, a combination of input, output, and computing hardware that can be used for work by an individual. More often, however, the term refers to a powerful stand-alone computer of the sort used in computer-aided design and other applications requiring a high-end,

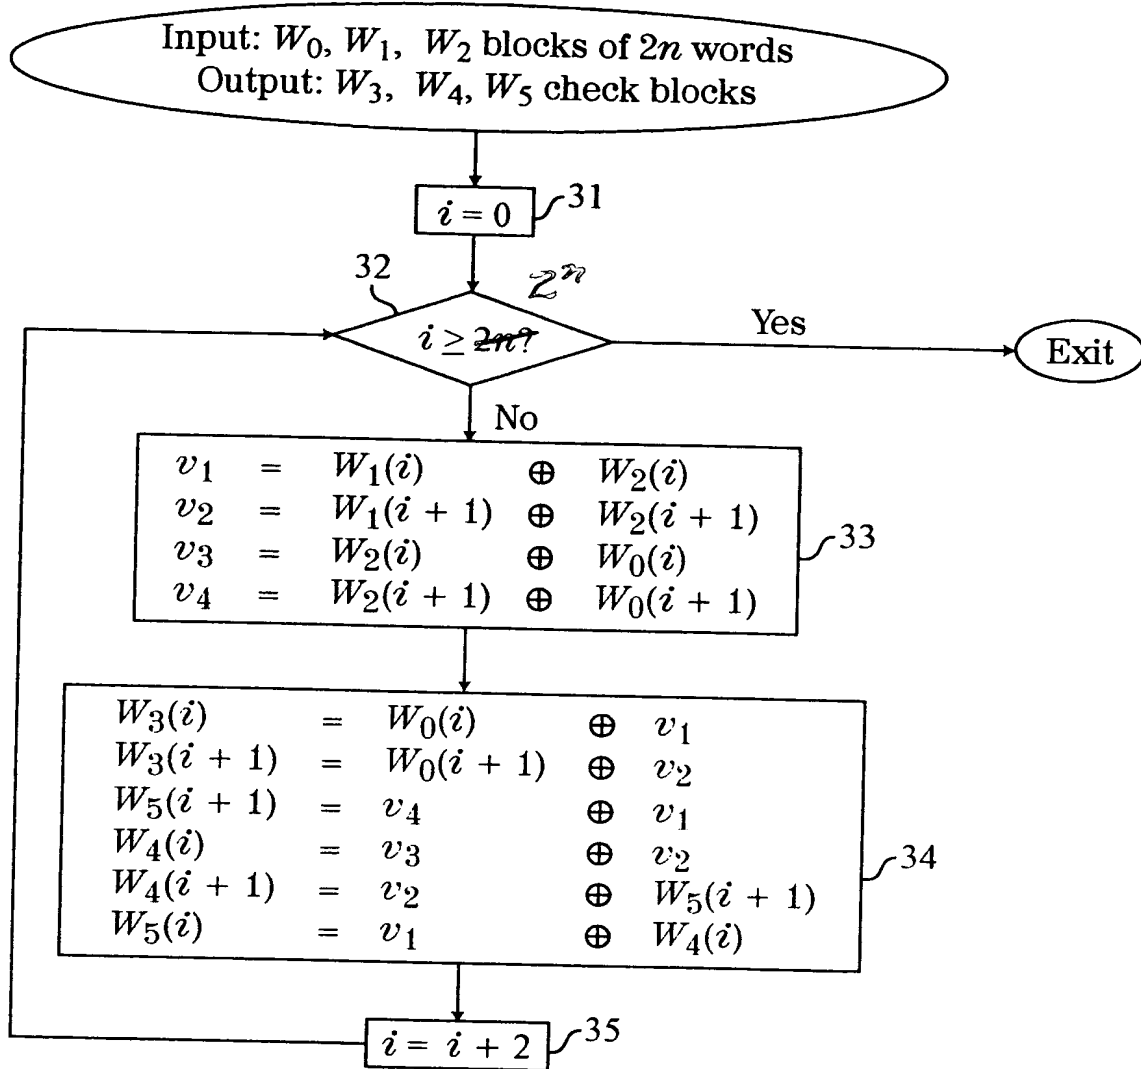


Figure 3

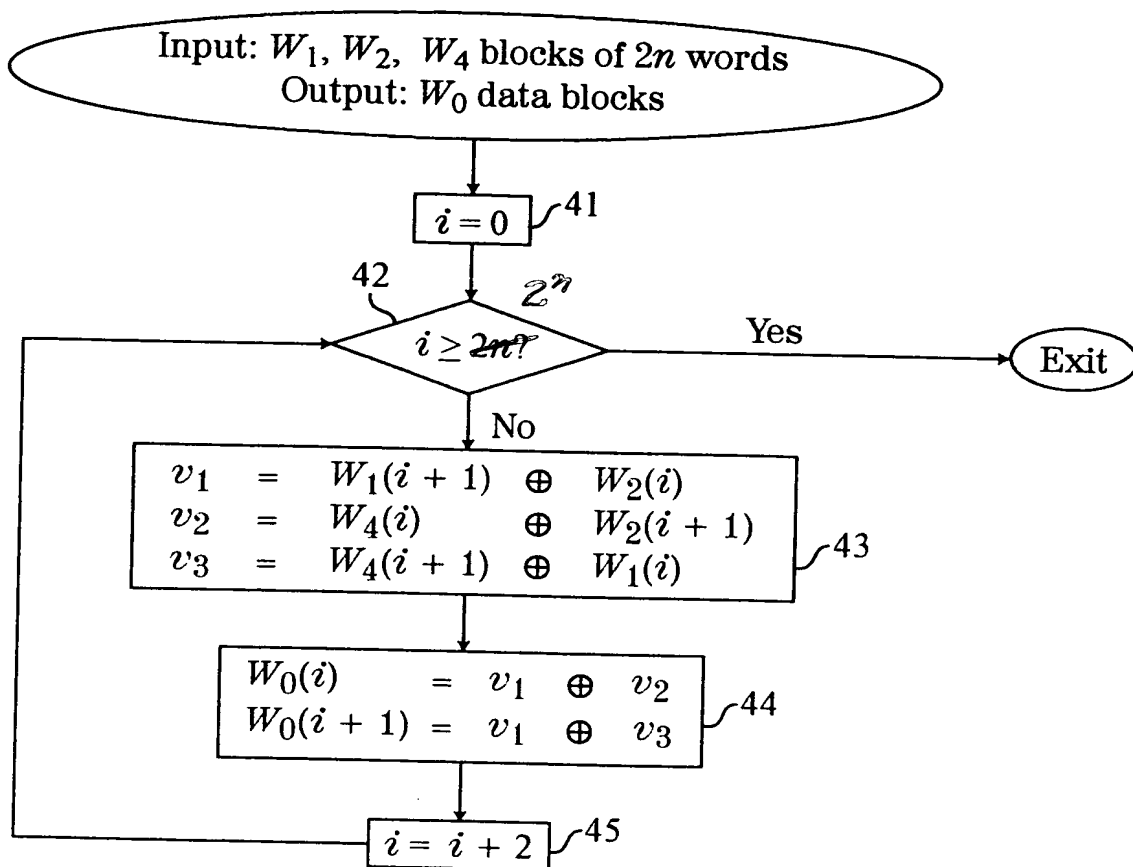


Figure 4